

# A tight bound on the speed-up through storage for quickest multi-commodity flows<sup>\*</sup>

Martin Groß<sup>1</sup> and Martin Skutella<sup>1</sup>

Fakultät II – Mathematik und Naturwissenschaften  
Institut für Mathematik, Sekr. MA 5-2  
Technische Universität Berlin, Straße des 17. Juni 136  
10623 Berlin, Germany

**Abstract.** Multi-commodity flows over time exhibit the non-intuitive property that letting flow wait can allow us to send flow faster overall. Fleischer and Skutella (IPCO 2002) show that the speed-up through storage is at most a factor of 2, and that there are instances where the speed-up is as large as a factor of  $4/3$ . We close this gap by presenting a family of instances for which the speed-up factor through storage converges to 2.

**Keywords:** dynamic flow, flow over time, quickest flow, holdover, multi-commodity, storage, waiting

## 1 Introduction

Multi-commodity flows over time are a useful tool for modelling problems that involve non-instantaneous transportation through some kind of network. Some recent examples for such problems include aircraft routing [1], cloud computing [2,3], road traffic networks [4], disaster management [5,6], evacuations [7], logistics [8,9], packet routing [10] and shared-ride systems [11].

During transportation from a source to a destination, depending on the particular application, making stops before reaching the destination might be feasible or not. For example, if the underlying network is a street network, we can only stop if parking space is available. Similarly, storage of data is possible in telecommunication networks only if buffers are present. Since keeping buffers or parking spaces available costs money, the question arises how much worse the situation becomes if we forbid storage compared to the setting where storage is allowed. We will refer to making stops during transportation as *storage* – this is sometimes referred to as *holdover* or *waiting* as well.

---

<sup>\*</sup> Supported by the DFG Priority Programme 1307 “Algorithm Engineering” (grant SK 58/6-3) and the DFG Research Center MATHEON “Mathematics for key technologies” in Berlin.

## 2 Quickest multi-commodity flows

We start by defining the *quickest multi-commodity flow problem*. For a more detailed introduction to the area of flows over time as well as for further pointers to the literature, we refer to [12,13].

We consider a network given by a directed graph  $G = (V, A)$  with a set of nodes  $V$  and a set of arcs  $A$ . Each arc  $a \in A$  has a capacity  $u_a \geq 0$  specifying the maximum rate at which flow may enter it at any point in time, and a transit time  $\tau_a \geq 0$  indicating how long flow needs to traverse from its start to its end node. Furthermore, we are given a set of  $k$  commodities  $K = \{0, 1, \dots, k-1\}$  with source nodes  $s_i \in V$ , sink nodes  $t_i \in V$ , and demands  $d_i \geq 0$  for every commodity  $i \in K$ .

We are interested in finding a quickest flow, that is, a flow over time that fulfills all supplies and demands as quickly as possible, i.e., within minimum time horizon  $T$ . More formally, a *multi-commodity flow over time*  $f$  with *time horizon*  $T > 0$  is a function  $f : (A \times K) \rightarrow ([0, T] \rightarrow \mathbb{R}_{\geq 0})$  that assigns a Lebesgue-integrable flow rate function  $f_a^i : [0, T] \rightarrow \mathbb{R}_{\geq 0}$  to every arc  $a \in A$  and every commodity  $i \in K$ , subject to *capacity* and *flow conservation* constraints that we specify below.

The value  $f_a^i(\xi)$  determines the *flow rate* of commodity  $i$  into arc  $a$  at time  $\xi$ . The flow rate specifies, as its name indicates, the rate at which the flow is entering the arc. This could for example be the number of cars entering a road per time unit at a given point in time. Due to the transit times, flow entering an arc  $a = (v, w)$  at time  $\xi$  will arrive at node  $w$  only at time  $\xi + \tau_a$ . While these flow rate functions are only defined on  $[0, T]$ , for the sake of convenience we set  $f_a^i(\theta) := 0$  for  $\theta \in \mathbb{R} \setminus [0, T]$ ,  $i \in K$ ,  $a \in A$ .

We can now specify the constraints on the flow rates. *Capacity constraints* make sure that the flow rates assigned to the commodities never violate the arc capacities:

$$\sum_{i \in K} f_a^i(\theta) \leq u_a \quad \text{for all } a \in A, \theta \in [0, T]. \quad (1)$$

*Flow conservation constraints* ensure that flow of a commodity can only emerge from its corresponding source node:

$$\sum_{a=(\cdot, v) \in A} \int_0^{\theta - \tau_a} f_a^i(\xi) d\xi - \sum_{a=(v, \cdot) \in A} \int_0^{\theta} f_a^i(\xi) d\xi \geq 0 \quad (2)$$

for all  $i \in K$ ,  $v \in V \setminus \{s_i\}$ ,  $\theta \in [0, T]$ .

Finally, in the end, at time  $T$ , all flow must have reached its corresponding sink node:

$$\sum_{a=(\cdot, v) \in A} \int_0^{T - \tau_a} f_a^i(\xi) d\xi - \sum_{a=(v, \cdot) \in A} \int_0^T f_a^i(\xi) d\xi = \begin{cases} d_i & \text{if } v = t_i, \\ -d_i & \text{if } v = s_i, \\ 0 & \text{otherwise,} \end{cases} \quad (3)$$

for all  $i \in K$ .

A feasible solution to the quickest multi-commodity flow problem is a flow over time  $f$  with time horizon  $T$  that satisfies constraints (1), (2), and (3). The objective is to minimize the time horizon  $T$ .

While for the case of a single commodity the quickest flow problem can be solved efficiently (cf. [14,15]), the multi-commodity case is weakly NP-hard [16]. On the positive side, there is a fully polynomial time approximation scheme (FPTAS) [17,18].


Notice that the flow conservation constraints (2) allow flow to arrive at an intermediate node  $v$ , wait there for some time, and continue onwards (i.e., the term on the left hand side of (2) might be positive at times). If storage of flow is prohibited at intermediate nodes, we can simply enforce equality in (2) for  $v \in V \setminus \{s_i, t_i\}$  at all times  $\theta$ . Flows over time obeying these *strict flow conservation* constraints are called *flows over time without (intermediate) storage*.

For the case of a single commodity, it follows from the work of Ford and Fulkerson [14] that there is always a quickest flow that does not make use of intermediate storage. On the other hand, Fleischer and Skutella [17,18] present an instance of the quickest multi-commodity flow problem where prohibiting intermediate storage increases the optimal time horizon by a factor of 4/3. On the positive side, they show that, for any instance, the optimal time horizon without storage is at most a factor of 2 larger than the optimal time horizon with storage. Since then it has been an open problem to close this gap (see, e.g., Exercise 10 in [12]). In the next section we present a family of instances for which the ratio between the optimal time horizon without and with storage, respectively, converges to 2. That is, the speed-up through storage can be as large as a factor of 2 and this bound is tight.

### 3 Instances with a large speed-up through storage

Before we give a formal definition of our family of instances, we first provide an intuitive description; see also Figure 1. Imagine a single lane roundabout traffic packed with cars where each car wants to make one full turn on the roundabout before leaving it. In an ideal world, these cars drive at uniform speed without interfering with one another and complete their turns simultaneously, then leave the roundabout after, shall we say, 10 seconds. Now we add another car to this scenario that is waiting to enter the roundabout at some point  $P$  in order to make a full turn as well. If the other cars are not willing to stop and wait, the new car has to wait until everyone else is done (as there are no gaps between cars) and only then take its turn to finish after 20 seconds. If however, every car is willing to shortly leave the roundabout at  $P$  in order to re-enter after only a short stop, the new car can join immediately and everyone is done after only slightly more than 10 seconds.

Let  $G = (V, A)$  be a directed cycle with  $k$  nodes  $V = \{v_0, v_1, \dots, v_{k-1}\}$  and arcs  $a_j = (v_j, v_{(j+1) \bmod k})$ ,  $j = 0, \dots, k-1$ . All arcs have unit capacities and transit times. There are  $k$  commodities  $K = \{0, 1, \dots, k-1\}$  with  $s_i = v_i$



**Fig. 1.** A single lane roundabout traffic (left) and the corresponding network  $G$  (right). All arcs have unit capacities and transit times.

and  $t_i = v_{(i-1) \bmod k}$ , for  $i = 0, \dots, k-1$ . Commodity 0 has demand  $d_0 := 2$  and all remaining commodities have unit demand  $d_i := 1$ ,  $i = 1, \dots, k-1$ .

**Lemma 1.** *If storage of flow at intermediate nodes is permitted, all demands can be sent within time horizon  $k+1$ .*

*Proof.* The following multi-commodity flow over time  $f$  sends all demands within time horizon  $k+1$ .

- For commodity 0, we send flow at rate 1 into the path  $s_0 = v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_{k-1} = t_0$  during the time interval  $[0, 2)$ . This flow does not wait at any node and thus arrives at its sink  $t_0$  before time  $2 + (k-1) = k+1$ .
- For commodities  $i = 1, \dots, k-1$ , we send flow at rate 1 into the path  $s_i = v_i \rightarrow \dots \rightarrow v_{k-1} \rightarrow v_0 \rightarrow \dots \rightarrow v_{i-1} = t_i$  during the time interval  $[0, 1)$ . For  $i > 1$ , flow of commodity  $i$  only waits at node  $v_0$  for one time unit before proceeding towards  $v_1$ . These flows thus arrive at their sinks before time  $1 + (k-1) + 1 = k+1$ .

By definition, the flow over time  $f$  sends all demands and satisfies flow conservation constraints. Thus, we only need to check the capacity constraints. By our choice of flow rates, capacity violations can only occur if two commodities send flow into an arc at the same time. This is prevented by the waiting of commodities  $2, \dots, k-1$  in  $v_0$ , thus completing the proof.

It remains to lower bound the time horizon necessary to send all demands without storage.

**Lemma 2.** *If storage of flow at intermediate nodes is forbidden, the minimal time horizon necessary to send all demands is larger than  $2k-2$ .*

*Proof.* We assume by contradiction that there is a flow over time  $f$  with time horizon  $T = 2k-2$  that routes all demands without storage at intermediate nodes. Since the flow paths of all commodities have transit time  $k-1$ , flow of

any commodity must leave its source before time  $k - 1$  and cannot reach its sink before time  $k - 1$ . In particular, flow cannot enter the last arc on its path before time  $k - 2$ . As a consequence, since flow may not wait at intermediate nodes and arc transit times are unit, all flow must be sent into some arc on its path within the time interval  $[k - 2, k - 1)$ , that is,

$$\sum_{j=0}^{k-1} \sum_{i \in K} \int_{k-2}^{k-1} f_{a_j}^i(\xi) d\xi \geq \sum_{i \in K} d_i = k + 1 .$$

On the other hand, by capacity constraints,  $\sum_{i \in K} f_{a_j}^i(\xi) \leq u_{a_j} = 1$  such that

$$\sum_{j=0}^{k-1} \sum_{i \in K} \int_{k-2}^{k-1} f_{a_j}^i(\xi) d\xi \leq k .$$

This contradiction concludes the proof.

Together, Lemma 1 and Lemma 2 yield the following theorem.

**Theorem 1.** *There is a family of quickest multi-commodity flow instances for which the ratio between the optimal time horizon without and with intermediate storage, respectively, converges to 2.*

We mention another interesting consequence of the family of instances discussed above. If we reduce the demand of commodity 0 to  $d_0 := 1$ , then there is a quickest flow with time horizon  $k$  which does not use storage at intermediate nodes. However, as soon as we increase  $d_0$  to  $1 + \varepsilon$  for any  $\varepsilon > 0$ , the same arguments as in the proof of Lemma 2 show that, if storage is prohibited, the optimal time horizon jumps above  $2k - 2$ . Thus, in contrast to the setting where storage is permitted (see Lemma 4.8 in [18]), the optimal time horizon is no longer a continuous function of the positive demand values if storage is prohibited.

## References

1. K. Roy, C. J. Tomlin, Solving the aircraft routing problem using network flow algorithms, in: Proceedings of the 2007 American Control Conference (ACC '07), 2007, pp. 3330–3335. doi:10.1109/ACC.2007.4282854.
2. A. Agapi, S. Soudan, M. Pasin, P. V.-B. Primet, T. Kielmann, Optimizing deadline-driven bulk data transfers in overlay networks, in: Proceedings of 18th International Conference on Computer Communications and Networks (ICCCN '09), 2009, pp. 1–8. doi:10.1109/ICCCN.2009.5235267.
3. S. Soudan, P. V.-B. Primet, Mixing malleable and rigid bandwidth requests for optimizing network provisioning, in: Proceedings of the 21st International Teletraffic Congress (ITC '09), 2009, pp. 1–8.
4. E. Köhler, R. H. Möhring, M. Skutella, Traffic networks and flows over time, in: J. Lerner, D. Wagner, K. A. Zweig (Eds.), Algorithmics of Large and Complex Networks: Design, Analysis, and Simulation, Vol. 5515 of Lecture Notes in Computer Science, Springer, 2009, pp. 166–196.

5. Y. Han, X. Guan, L. Shi, Optimization based method for supply location selection and routing in large-scale emergency material delivery, *IEEE Transactions on Automation Science and Engineering* 8 (2011) 683–693. doi:10.1109/TASE.2011.2159838.
6. M. S. Osman, B. Ram, J. Bhadury, P. Stanfield, L. Davis, F. Samanlioglu, Optimization model for distributed routing for disaster area logistics, in: *Proceedings of the 5th IEEE International Conference on Service Operations, Logistics and Informatics (SOLI '09)*, 2009, pp. 278–283. doi:10.1109/SOLI.2009.5203944.
7. M. Lahmar, T. Assavapokee, S. A. Ardekani, A dynamic transportation planning support system for hurricane evacuation, in: *Proceedings of the 9th IEEE Intelligent Transportation Systems Conference (ITSC '06)*, 2006, pp. 612–617. doi:10.1109/ITSC.2006.1706809.
8. G. Andreatta, L. De Giovanni, G. Salamaso, Fleet quickest routing on grids: A polynomial algorithm, *International Journal of Pure and Applied Mathematics* 62 (2010) 419–432.
9. S. Hernández, S. Peeta, G. Kalafatas, A less-than-truckload carrier collaboration planning problem under dynamic capacities, *Transportation Research Part E: Logistics and Transportation Review* 47 (2011) 933–946. doi:10.1016/j.tre.2011.03.001.
10. B. Peis, A. Wiese, Universal packet routing with arbitrary bandwidths and transit times, in: O. Günlük, G. J. Woeginger (Eds.), *Proceedings of the 14th Conference on Integer Programming and Combinatorial Optimization (IPCO '10)*, Vol. 6655 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 362–375. doi:10.1007/978-3-642-20807-2\_29.
11. M. Braun, S. Winter, Ad-hoc solution of the multicommodity-flow-over-time problem, *IEEE Transactions on Intelligent Transportation Systems* 10 (2009) 658–667. doi:10.1109/TITS.2009.2026443.
12. M. Skutella, An introduction to network flows over time, in: *Research Trends in Combinatorial Optimization*, Springer, 2009, pp. 451–482. doi:10.1007/978-3-540-76796-1\_21.
13. M. Groß, Approximation algorithms for complex network flow over time problems, Ph.D. thesis, TU Berlin (2014).
14. L. R. Ford, D. R. Fulkerson, Constructing maximal dynamic flows from static flows, *Operations Research* 6 (1958) 419–433.
15. R. E. Burkard, K. Dlaska, B. Klinz, The quickest flow problem, *Zeitschrift für Operations Research – Methods and Models of Operations Research* 37 (1993) 31–58.
16. A. Hall, S. Hippler, M. Skutella, Multicommodity flows over time: Efficient algorithms and complexity, *Theoretical Computer Science* 379 (2007) 387–404.
17. L. Fleischer, M. Skutella, The quickest multicommodity flow problem, in: W. J. Cook, A. S. Schulz (Eds.), *Proceedings of the 9th Conference on Integer Programming and Combinatorial Optimization (IPCO '02)*, Vol. 2337 of *Lecture Notes in Computer Science*, Springer, 2002, pp. 36–53. doi:10.1007/3-540-47867-1\_4.
18. L. Fleischer, M. Skutella, Quickest flows over time, *SIAM Journal on Computing* 36 (2007) 1600–1630.